

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 5 : G06F 15/18		A1	(11) International Publication Number: WO 90/11568 (43) International Publication Date: 4 October 1990 (04.10.90)
(21) International Application Number: PCT/US90/00828 (22) International Filing Date: 21 February 1990 (21.02.90)		(74) Agent: BRUNS, Gregory, A.; Honeywell Inc., MN12-8251, Honeywell Plaza, Minneapolis, MN 55408 (US).	
(30) Priority data: 329,623 28 March 1989 (28.03.89) US		(81) Designated States: AT (European patent), BE (European patent), CA, CH (European patent), DE (European patent), DK (European patent), ES (European patent), FR (European patent), GB (European patent), IT (European patent), JP, LU (European patent), NL (European patent), SE (European patent).	
(71) Applicant: HONEYWELL, INC. [US/US]; Honeywell Plaza, Minneapolis, MN 55408 (US).		Published <i>With international search report.</i>	
(72) Inventors: GUHA, Aloke ; 5104 10th Avenue South, Minneapolis, MN 55417 (US). HARP, Steven, A. ; 2245 Scudder, St. Paul, MN 55108 (US). SAMAD, Tariq ; 4949 Russell Avenue South, Minneapolis, MN 55410 (US).			
(54) Title: GENETIC SYNTHESIS OF NEURAL NETWORKS			
(57) Abstract			
<p>The invention relates to a method for using genetic type learning techniques in connection with designing a variety of neural networks that are optimized for specific applications. In the invention herein a general representation of neural network architectures is linked with the genetic learning strategy to create a flexible environment for the design of custom neural networks. A concept upon which the invention is based is the representation of a network design as a "genetic blueprint" wherein the recombination or mutation of subsequently generated editions of such blueprints result in different but related network architectures.</p>			

* See back of page

DESIGNATIONS OF "DE"

Until further notice, any designation of "DE" in any international application whose international filing date is prior to October 3, 1990, shall have effect in the territory of the Federal Republic of Germany with the exception of the territory of the former German Democratic Republic.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	ES	Spain	MG	Madagascar
AU	Australia	FI	Finland	ML	Mali
BB	Barbados	FR	France	MR	Mauritania
BE	Belgium	GA	Gabon	MW	Malawi
BF	Burkina Faso	GB	United Kingdom	NL	Netherlands
BG	Bulgaria	HU	Hungary	NO	Norway
BJ	Benin	IT	Italy	RO	Romania
BR	Brazil	JP	Japan	SD	Sudan
CA	Canada	KP	Democratic People's Republic of Korea	SE	Sweden
CF	Central African Republic	KR	Republic of Korea	SN	Senegal
CG	Congo	LI	Liechtenstein	SU	Soviet Union
CH	Switzerland	LK	Sri Lanka	TD	Chad
CM	Cameroon	LU	Luxembourg	TG	Togo
DE	Germany, Federal Republic of	MC	Monaco	US	United States of America
DK	Denmark				

GENETIC SYNTHESIS OF NEURAL NETWORKS

The invention hereof relates to a method for using genetic type learning techniques in connection with designing a variety of neural networks that are optimized for specific applications.

Previous work in the design of neural networks has revealed the difficulty in determining an appropriate network structure and good values for the parameters of the learning rules for specific applications.

The genetic algorithm is an optimization method based on statistical selection and recombination. The method is inspired by natural selection. A few researchers (Dolan & Dyer (1987), Dress & Knisely (1987) Davis (1988), Montana and Davis (1989) and Whitley (1988)) have applied generic algorithms in a limited fashion to generate neural networks for specific problems. For example, Davis and Montana (1988, 1989) and Whitley (1988) use the genetic algorithm to adjust weights given a fixed network structure.

- 2 -

In the invention herein a general representation of neural network architectures is linked with the genetic learning strategy to create a flexible environment for the design of custom neural networks. A 5 concept upon which the invention is based is the representation of a network design as a "genetic blueprint" wherein the recombination or mutation of subsequently generated editions of such blueprints result in different but related network architectures.

10 To illustrate the invention there is described herein a system for the genetic synthesis of a particular class of neural networks that we have implemented. Our current implementation is restricted to network structures without feedback connections and 15 incorporates the back propagation learning rule. The invention can, however, be used for arbitrary network models and learning rules.

20 The method herein involves the use of genetic algorithm methods to design new neural networks. The genetic algorithm (GA) is a robust function optimization method. Its use is indicated over gradient descent techniques for problems fraught with local minima, discontinuity, noise, or large numbers of dimensions. A 25 useful feature of the GA is that it scales extremely well, increasing dimensionality has comparatively little

- 3 -

effect on performance. The first step in the application of the GA to a function is the encoding of the parameter space as a string of (typically binary) digits. Substrings in such a representation correspond to parameters of the function being optimized. A particular individual bit string (i.e. some choice of 1 or 0 for each position) represents a point in the parameter space of the function. The GA considers a population of such individuals. The population, in conjunction with the value of the function for each individual (generally referred to as "fitness"), represents the state of the search for the optimal string. The GA progresses by implicitly encoding information about the function in the statistics of the population and using that information to create new individuals. The population is cyclically renewed according to a reproductive plan. Each new "generation" of the population is created by first sampling the previous generation according to fitness; the method used for differential selection is known to be a near-optimal method of sampling the search space. Novel strings are created by altering selected individuals with genetic operators. Prominent among these is the crossover operator which synthesizes new strings by splicing together segments of two sampled individuals.

- 4 -

A main object of the invention is to provide a new method as referred to above for designing optimized artificial neural networks.

Other objects and advantages of the invention will become apparent from the following specification, appended claims and attached drawings.

In the drawings:

Fig. 1 illustrates a multilayer neural network of the type which may be designed for a specific purpose in accordance with the method of the present invention;

Fig. 2 illustrates schematically how a population of "blueprints" (designs for different neural networks) is cyclically updated by the genetic algorithm based on their fitness;

Fig. 3 shows schematically an example of a three-layer network which may be described by a bit string representation in accordance with the invention;

Fig. 4 illustrates a bit string representation which facilitates practicing the invention;

Fig. 5 illustrates the gross anatomy of a network representation having areas or layers 0 to N;

Fig. 6 illustrates an arrangement of areas (or layers) and projections extending therebetween;

Fig. 7 shows the spatial organization of layers;

-5-

Figs. 8 and 9 show examples of absolute and relative addressing for specifying the target destinations of projections which extend from one layer to another layer;

5 Figs. 10 to 12 show illustrative examples of the area specification substring of Fig. 4;

Figs. 13a, 13b and 13c show projection features relating to connections between layers of the network;

10 Figs. 14a, 14b and 14c show a schematic example of a specific network structure generated by the method herein at different levels of detail;

Fig. 15 shows the basic reproductive plan used in experiments pursuant to the method herein;

15 Figs. 16a, 16b and 16c show an example of the operation of a genetic operator;

Fig. 17 shows the principle data structures in a current implementation program with one individual being shown parsed and instantiated; and

20 Figures 18 to 21 show performance curves relating to the rate of learning of networks.

The method herein relates to the designing of multilayer artificial neural networks of the general type 10 shown in Fig. 1. The network 10 is illustrated as having three layers (or areas) 12, 14 and 16 but could have more than three layers or as few as one layer if desired. Each of the layers has computational units

SUBSTITUTE SHEET

- 6 -

18 joined by connections 19 which have variable weights associated therewith in accordance with the teaching of the prior art.

5 In this and other figure connections are shown in the forwardly feeding direction. The invention is not limited to this construction, however, and feedback connections may also be accommodated, for example.

10 Also, the scope of the network design method disclosed herein is not limited to the design of the network shown in Fig. 1.

15 Fig. 2 illustrates schematically how a population of blueprints 20 (i.e. bit string designs for different neural networks) are cyclically updated by a genetic algorithm based on their fitness. The fitness of a network is a combined measure of its worth on the problem, which may taken into account learning speed, accuracy and cost factors such as the size and complexity of the networks.

20 The method begins with a population of randomly generated bit strings 20. The actual number of such bit strings is somewhat arbitrary but a population size of 30 to 100 seems empirically to be a good compromise between computational load, learning rate and genetic drift.

- 7 -

NEURAL NETWORK LEARNING ALGORITHMS

Learning approaches for neural networks fall into three general categories: unsupervised learning, reinforcement learning, and supervised learning. In unsupervised learning, the network receives no evaluative feedback from the environment; instead it develops internal models based on properties of received inputs. In reinforcement learning, the environment provides a weak evaluation signal. In supervised learning the "desired output" for the network is provided along with every training input. Supervised learning, specifically back propagation, is used to illustrate the invention but in concept the invention can be used with any learning approach.

The set of input-output examples that is used for supervised learning is referred to as the training set. The learning algorithm can be outlined as follows: FOR EACH (training-input, desired-output) pair in the Training-Set

- 20
 - o Apply the training-input to the input of the network.
 - o Calculate the output of the network.
 - o IF the output of the network ≠ desired-output
 - 25 o THEN modify network weights

- 8 -

The entire loop through the training set, referred to as an epoch, is executed repeatedly. One or both of two termination criteria are usually used: there can be a lower bound on the error over an epoch and/or a 5 limit on the number of epochs. Training a network in this fashion is often very time consuming. Until better learning techniques become available, it is best to plan the training phase as an "off-line" activity. Once trained, the network can be put to use. The 10 computational demands of such a network during the operational phase can usually be satisfied with only rudimentary hardware for many interesting applications.

The neural network learning approach which we have currently implemented is the well-known 15 backpropagation algorithm. (Werbos, 1974; Le Cun, 1986; Parker, 1985; Rumelhart, Hinton & Williams, 1985).

The backpropagation algorithm is described in Appendix B.

20 BLUEPRINT REPRESENTATIONS

The invention herein is mainly directed to a representation of the blueprint 20 that specifies both the structure and the learning rule, the genetic algorithm parameters that determine how the genetic 25 operators are used to construct meaningful and useful

- 9 -

network structures, and the evaluation function that determines the fitness of a network for a specific application.

The representation of blueprints or bit strings 20 for specialized neural networks should ideally be able to capture all potentially "interesting" networks, i.e., those capable of doing useful work, while excluding flawed or meaningless network structures. It is obviously advantageous to define the smallest possible search space of network architectures that is sure to include the best solution to a given problem. 20 An important implication of this goal in the context of the genetic algorithm is that the representation scheme should be closed under the genetic operators. In other words, the recombination or mutation of network blueprints should always yield new, meaningful network blueprints. There is a difficult trade off between expressive power and the admission of flawed or 25 uninteresting structures.

- 10 -

Fig. 3 shows schematically an example of how each layer of a three-layer network may be described in accordance with the invention by a bit string representation which comprises three substrings 17. The format for a single substring 17 is shown in more detail in Fig. 4.

The gross anatomy of a multilayer network representation 20 having substring layers or areas 17 (Area 0 to Area N) is illustrated in Figure 5.

Conceptually, all of the parameters for a single network are encoded in one long string of bits which is the representation 20 of Figure 5. The patterned bars are markers indicating the start and end of the individual area or layer segments 17.

The term projection as used herein has reference to the grouping or organization of the connections 19 which extend between the computational units 18 of the layers of the networks such as in the network illustrations of Figs. 1 and 3.

In Fig. 1 the input connections to layer 12 represent a single input projection and the output connections extending outwardly from layer 16 represent a single output projection. Likewise the connections 19 between layers 12 and 14 represent a single projection and the connections 19 between the layers 14 and 16 represent a single projection.

- 11 -

An example of a projection arrangement for a particular network is shown in Fig. 6 with projections 22 to 28 being illustrated for layers or areas 31 to 35. Of interest is that layer 32 has two projections 24 and 25 extending respectively to layers 33 and 34. Also of interest is the opposite arrangement wherein layer 35 receives projections 26 and 27 from layers 33 and 34 respectively.

5

Each of the projections is represented by three lines which signify that each projection consists of a predetermined or desired plurality of only the connections 19 which extend between two particular layers.

10

Referring to Fig. 4, it will be apparent that an area or layer specification substring 17 as illustrated in this figure is applicable to each one of the layers 12, 14 and 16 of the network 10 of Fig. 1.

15

A bit string 20 is thus composed of one or more segments or substrings 17, each of which represents a layer or area and its efferent connectivity or projections. Each segment is an area specification substring 17 which consists of two parts:

20

- o An area parameter specification (APS) which is of fixed length, and

25

- 12 -

parameterizes the area or layer in terms of its address, the number of units in it, and how they are organized.

- o One or more projection specification fields (PSFs), each of fixed length. Each such field describes a connection from one layer to another layer. As the number of layers is not fixed in this architecture (although bounded), the length of this field will increase with the number of projection specifiers required. A projection (e.g., on the projections 22 to 28 in Fig. 6) is specified by the address of the target area, the degree of connectivity and the dimension of the projection to the area, etc.

The fact that there may be any number of areas 17 motivates the use of markers with the bit string to designate the start and end of APSSs and the start of 20 PSFs. The markers enable a reader program to parse any well-formed string into a meaningful neural network architecture. The same markers also allow a special genetic crossover operator to discover new networks without generating "nonsense strings". Markers are 25 considered "meta-structure": they serve as a framework but don't actually occupy any bits.

- 13 -

Fig. 4 shows how the APS and PSF are structured in our current representation. The portions of the bit string representing individual parameters are labeled boxes in the figure. They are substrings consisting of some fixed number of bits. Parameters described by an interval scale (e.g. 0, 1, 2, 3, 4) are rendered using Gray coding, thus allowing values that are close on the underlying scale to be close in the bit string representation (Bethke, 1980, Caruana & Schaffer, 1988).

5

In the APS, each area or layer has an identification number that serves as a name. The name need not be unique among the areas of a bit string. The input and output areas have the fixed identifiers, 0 and 7 in the embodiment herein.

10

An area also has a size and a spatial organization. The "total size" parameter determines how many computational units the area will have. It ranges from 0 to 7, and is interpreted as the logarithm (base 2) of the actual number of units; e.g., if total size is 5, there are 32 units. The three "dimension share" parameters, which are also base 2 logarithms, impose a spatial organization on the units. The units of areas may have 1, 2 or 3 dimensional rectilinear extent, as illustrated in Fig. 7.

15

20

25

- 14 -

The motivation for this organization comes from the sort of perceptual problems to which neural networks are apparently well suited. For example, an image processing problem may best be served by square array,
5 while an acoustic interpretation problem might call for vectors. The organization of the units in more conventional approaches is often left implicit. In the invention herein dimensionality has definite implications for the architecture of projections such as
10 the projections 22 to 28 of Fig. 6.

The PSFs in an area's segment of the bit string determine where the outputs of units in that layer will (attempt to) make efferent connections, and how. The representation scheme does not assume a simple pipeline
15 architecture, as is common. Fig. 6, for example, shows a five-area network in which projections split from the second area and rejoin in the fifth.

Each PSF indicates the identity of the target area. There are currently two ways it can do this, distinguished by the value of a binary addressing mode parameter in each PSF. In the "Absolute" mode, the PSF's address parameter is taken to be the ID number of the target area. Some examples of absolute addressing are shown in Fig. 8.

- 15 -

5 The "Relative" mode indicates that the address bits hold the position of the target area in the bit string relative to the current area. A relative address of zero refers to the area immediately following the one containing the projection; a relative address of n refers to the nth area beyond this, if it exists.

10 Relative addresses indicating areas beyond the end of the blueprint are taken to refer to the final area of the blueprint—the output area. Some examples of relative addressing are shown in Fig. 9.

15 The purpose of different addressing schemes is to allow relationships between areas to develop, and be sustained and generalized across generations through the genetic algorithm's reproductive plan. Specifically, the addressing schemes are designed to help allow these relationships to survive the crossover operator, either intact or with potentially useful modifications.

20 Absolute addressing allows a projection to indicate a target no matter where that target winds up in the chromosome of a new individual. Relative addressing helps areas that are close in the bit string to maintain projections, even if their IDs change.

25 Referring to Figs. 10 to 12, the dimension radii parameters (also base 2 logarithms) allow units in an area to project only to a localized group of units in

- 16 -

the target area. This feature allows the target units to have localized receptive fields 29, which are both common in biological neural networks and highly desirable from a hardware implementation perspective.

5 Even within receptive fields 29, projections between one area or layer and another do not necessarily imply full factorial connectivity. The connection density parameter for the projection may stipulate one of eight degrees of connectivity between 30% and 100%.

10 At this point it may be well to mention that, because of the magnitude of the numbers involved for the units 18 and the connections 19, it is contemplated that in a typical system the numbers will be represented by their logarithms. In Figs. 10 to 12 and 15 herein which 15 show examples of the substring 17, decoded numbers are used by way of illustration to facilitate an understanding of the concepts.

20 Projections include a set of weighted connections. The weights are adjusted by a learning rule during the training of the network. Parameters are included in the PSF to control the learning rule for adjusting the weights of the projection. The eta parameter controls the learning rate in back propagation and may take on one of 8 values between 0.1 and 12.8.

25 Eta need not remain constant throughout training. A

- 17 -

separate eta-slope parameter controls the rate of exponential decay for eta as a function of the training epoch.

5 An example of how this representation scheme can be used to specify a 3-layer network is shown in Fig. 3.

10 The first and last areas or layers of the network have a special status. The first, the input area, represents the set of terminals that will be "clamped" by the network's environment, effectively the input stimulus. The final area is always the output area, and has no projections.

15 A blueprint representation in BNF of the neural network described herein is Appendix A at the end hereof. It is anticipated that there will be future modifications and additions to it.

20 Figs. 10 to 12 show three examples of substrings 17 which illustrate the projection specifier sections thereof relative to the radius and the connection density parameters. These figures show examples of projections 21 from a layer or Area 1 to a layer or Area 2. The projection in Fig. 10 is from a one dimensional area (Area 1) to a two dimensional area (Area 2) and the projections in Figs. 11 and 12 are each from a one dimensional area (Area 1) to another dimensional area (Area 2).

25

- 18 -

In Fig. 10 the illustrated projection is to an 8 by 4 projection array 29 of computational units 18 and, by convention, this array is deemed to have a radius of 4 in the vertical direction and a radius of 2 in the horizontal direction. The object array 29 is 5 symmetrically arranged relative to the source unit 18a in Area 1. As each unit within the boundary of the projection array 19 is connected, the connection density parameter is 100.

10 It will be understood that each of the computational units 18 in Area 1 will in a similar manner have connections to respective 8 x 4 projection arrays of units in Area 2 which results in substantial overlapping of projection arrays and a very dense 15 connection system.

In Fig. 11 the projections are to every other one of a linear array of 20 units. The radius is 8 indicated but the connection density parameter is only 20 50 because only half of the units within the radius are connected.

Fig. 12 is similar to Fig. 11 except that every computational unit in the array is connected and thus the connection density is 100.

Figs. 11 and 12 are similar to Fig. 10 relative 25 to the matter of having each unit in Area 1 connected to a respective projection array of units in Area 2.

-19-

Potential target units of a projection from a given source unit are determined by radii along three dimensions. Figures 13a, 13b and 13c are three two-dimensional examples of this.

5 Figures 14a, 14b and 14c, taken together, provide a schematic example of a specific network structure generated by the method herein.

ADAPTING GENETIC ALGORITHMS

10 The version of the genetic algorithm used in the method herein employs a reproductive plan similar to that described by Holland (1975) as "type R". The basic plan for generating each new generation is given in Fig. 15. The sampling
15 algorithm is based on the stochastic universal sampling scheme of Baker (1987). This is preferred for its efficiency and lack of bias. Some of the details are not shown by the diagram. A final step was added to insure that the best individual from
20 generation i was always retained in generation i+1.

The genetic algorithm (GA) itself has a number of parameters. Good values for these are important to the efficient operation of the system. These parameters include the population size, the
25 rates at which to apply the various genetic operators, and other aspects of the synthetic ecology.

SUBSTITUTE SHEET

- 20 -

Two genetic operators have been used: crossover and mutation. The crossover operator effectively exchanges homologous segments from the blueprints of two networks from the current generation to create a blueprint for a network in the next generation. In most applications of the genetic algorithm, homologous segments are identifiable by absolute positions in the bit string. For example, The Nth bit will always be used to specify the same trait in any individual. Because the representation herein allows variable length strings, a modified two-point crossover operator was employed that determined homologous loci on two individuals by referring to the string's markers, discussed above. The decision to use a two-point crossover as opposed to the more common single-point version was motivated by Booker's (1987) report that improved off-line performance could be achieved this way.

The mutation operator was used at a low rate to introduce or reintroduce alleles-alternate forms of the same functional gene. Current applications of the genetic algorithm have demonstrated an effective contribution from mutation at rates on the order of 10^{-2} or less.

-21-

Despite the fact that the bit string representation was designed with closure under the genetic operators in mind, it is still possible for the GA to generate individuals that are *prima facie* unacceptable. A blatant example would be a network plant that had no pathway of projections from input to output. Subtler problems arise from the limitations of our simulation capability. In our initial work we have limited recurrence; network plans with feedback cannot be tolerated under simple back propagation. Two strategies have been employed for minimizing the burden of these misfits. First, the reproductive plan culls individuals with fatal abnormalities; individuals with fatal abnormalities; individuals with no path from input to output area compose the bulk of this group. Second, blueprints with minor abnormalities are "purified" in their network implementation, i.e. their defects are excised.

Figures 16a, 16b and 16c show an example of how the crossover operator can create new strings with different values for fields than either of the parents. Here it is assumed that the fields use a simple binary encoding scheme.

- 22 -

EVALUATION OF SYNTHESIZED NETWORKS

Suitable improvements over generations can only be accomplished if the evaluation function used to measure the fitness of a network is appropriate. A measure of fitness is necessary for the GA to produce better and better networks. It is helpful to envision the algorithm as exploring the surface over the blueprint representation space defined by this function in an attempt to locate the highest peaks.

10 In accordance with the requirements of the evaluation function stated above, we have initially formulated the evaluation function as a weighted sum of the performance metrics, p_i . The evaluation function, $F(i)$, for individual i can be expressed as:

15

$$F(i) = \sum_{j=1}^n a_j p_j (i)$$

20 The coefficients a_j may be adjusted by the user to reflect the desired character of the network. Metrics that have been considered thus far include performance factors such as observed learning speed and the performance of the network on noisy inputs, and cost factors such as the size of the network, and the number of connections formed. We have adopted a melange of 25 different performance and cost factors since performance

- 23 -

criteria vary from application to application. Because the relative weight on each factor can be modified, the network structure can be tuned for different optimization criteria. For example, if one of our goals is to synthesize networks that are computationally efficient, the size metrics might be given negative weights. On the other hand, if accuracy and noise tolerance is more crucial, then the performance on noisy input patterns would be given a higher weight.

10

EVALUATION OF GA PERFORMANCE

In order to make conclusions about the performance of the genetic algorithm (as opposed to the networks themselves) in discovering useful architectures, we require some standard to compare it against. This is difficult since there seems to be no published data directly relevant to the problem. Our approach is to run a control study in which network blueprints are generated at random, evaluated, and the best retained. This is effected by simply "turning off" the genetic operators of crossover and mutation. Random search is an oft employed benchmark of performance that other search algorithms must exceed to demonstrate their value.

20

25

- 24 -

DATA STRUCTURES

The major data structures in a current implementation of the invention are objects that are created and linked together at run time. The most 5 prominent object is the "experiment" which maintains the current population, the history of performance over generations, as well as various control and interface parameters. The performance history is a list of records, one per generation, noting, among other things, 10 on-line, off-line, average and best scores. The population comprises the individuals of the current generation as shown in Fig. 17.

Each individual has an associated blueprint, which is stored as a bundle of bit vectors [bit vectors 15 are one-dimensional arrays in which each element occupies one bit in the machine's memory].

The bit vectors are of two types, areas (APS) and projections (PSF), as indicated by the BNF. The structure of each type is defined by a Lisp form, 20 indicating the names of each field, and how many bits it should occupy. For example, the projection specification is defined as:

```
(def-bit-vector PROJECTION-SPEC
  (radius-1 3)
  (radius-2 3)
  (radius-3 3)
  (connection-density 3)
  (target-address 3)
  (address-mode 1)
  (initial-eta 3)
  (eta-slope 3) )
```

- 25 -

This form automatically defines the accessors needed to extract the value for each parameter from any given bit vector. The accessors transparently effect the gray coding and decoding fields. Most of the integral values of fields are interpreted through lookup tables; for example, an eta table translates the values 0...7 to etas from 0.1 to 12.8.

5

Genetic operators such as crossover and mutation directly modify this bit vector blueprint, which is considered the master plan for the individual. Pieces of it are actually shared with its offspring. The bit vectors are not directly useful in running an actual neural network, however. For this, the individual must be parsed, purified, and instantiated.

10

15 When an individual is parsed, the bit string form of the blueprint is translated into a network of nodes—an area node for each area, and a projection node for each projection. Parsing works out the inter-area addressing done by projections, and the nodes carry

20

parameter values interpreted from the associated bit vectors. The network, or parsed blueprint, is associated with the object representing individual.

25

A parsed blueprint may have defects that prevent a meaningful interpretation as a neural network. For example, it might contain projections with no valid target, or projections indicating feedback

- 26 -

circuits, which are prohibited in the current implementation. Rather than discarding slightly imperfect individuals, an attempt is made to patch them after parsing. The patching step is called 5 purification. The purifier removes dangling nodes and cuts circuits in an attempt to create a viable individual while making as few changes as possible.

Following parsing and purification, an individual is instantiated. Instantiation involves 10 allocating and initializing vectors for units, weight matrices, mask matrices, threshold vectors, and other numerical storage. References to these data objects are kept in the nodes of the individual's parsed blueprint.

15 THE EVALUATION PROCESS

The purpose of the parse/purify/instantiate sequence is to set the stage for the evaluation of the individual, i.e. the computation of a score. The score 20 is a weighted sum of a set of performance metrics. The weights may be set by the user at run time.

Some of these metrics are immediate consequences of instantiation, e.g. number of weights, number of units, number of areas, and average fan-out. Other metrics depend on the individual network's 25 performance on a given problem (such as digit

- 27 -

recognition). Examples of such metrics are: the learning rate of the network, its final performance on the training set, its performance on non-degraded inputs and on novel inputs, and its performance after 5 temporarily nullifying a random sample of either the weights or units of the network.

RESULTS, ANALYSIS AND DISCUSSION

Despite the restricted scope of initial 10 experiments, the method herein has produced reasonable networks, and has achieved significant improvements over the chance structures in its initial generation. In most cases, the networks produced have been structurally fairly simple.

15

PERFORMANCE CRITERIA

There are several common ways to look at the changes in population performance over time in genetic optimization systems, and most of our charts include 20 four. Because our reproductive plan goes through separate phases of reproduction and evaluation, the data points are actually recorded at the end of each generation.

25

- 28 -

Define S_i to be the score of the i th individual generated. The best function indicates the performance of the best individual discovered by the GA up to a given time, i.e.

5

$$\text{Best}(i) = \text{Max } [S_j, j=1, \dots, i]$$

The off-line GA performance is the mean of the best

individual's scores found up to a give time:

10

$$\text{Off-line}(i) = \frac{1}{i} \sum_{j=1}^i \text{Best}(j)$$

An alternative is the on-line performance. This is simply the mean of all individuals scores evaluated so far. At the end of time i , this. would be:

15

$$\text{On-line}(i) = \frac{1}{i} \sum_{j=1}^i S_j$$

20

Another interesting function is the average score for all of the individuals in a given generation. If G_i is the set of individuals in the i th generation, then:

25

$$\text{Average}(i) = \frac{1}{|G_i|} \sum_{j \in G_i} S_j$$

- 29 -

On-line performance is perhaps most relevant to systems that must interact with a real-time process, whereas off-line performance is more relevant to systems that are concerned only with finding the best and not how much it costs to look. For example, if one were picking horses, it would be important to take into consideration all of the poor bets as well as the winners, motivating interest in on-line performance. If one were optimizing a function, the only concern might be about the quality of the best point tested, motivating off-line performance. Noting that the "Best" and "Offline" functions are isotone by definition they can only increase or remain constant over the course of an experiment, and cannot decrease.

15

EXPERIMENT 1

Application: Digit Recognition

20

Optimization Criterion: Area under learning curve

Population Size: 30

Generations: 60

25

The average performance of the network population increased eight-fold from the first to the sixtieth generation. The network learned to criterion in 48 epochs.

- 30 -

Since only one factor was being directly optimized, others such as the number of weights were free to vary. The network had 1481 weights. A network which had exactly one weight connecting each input with each output would have only a third as many weights. Such networks were also produced, and these learned perfectly as well, but took more than twice as long. The performance of this experiment for this 60 generation experiment is summarized by Fig. 18.

In the initial generations, hidden-layer structures were present. It was not obvious to us that this problem is linearly separable until the experiment started producing two-layer structures that were learning perfectly. Since hidden layers are not needed for this problem, and since learning rates in general degrade as hidden layers are added to a network (although this degradation is much less severe with the modified back-propagation rule we are using [Samad, 1988] than with the original rule), towards the end of the simulation multiple-layer structures were rare.

In order to evaluate the performance of the GA in discovering better networks, the digit recognition problem was repeated with the GA disabled. To achieve this, random individuals were generated where crossover or mutation would have been applied. Again, scores were

- 31 -

based exclusively on the area under the learning curve. The results of this experiment are charted in Fig. 19.

5 While the random search and GA experiments started with a very similar populations in generation 0, the performance of the two algorithms soon diverged. In particular, average and on-line performances of the random search algorithm were conspicuously inferior to the GA. This is to be expected if the GA is successful in retaining some of the better characteristics from one 10 generation to the next; the random search procedure is confined to picking "losers" at a fairly constant rate. The off-line performance is arguably a more interesting comparison for this problem between the GA and random search. Figure 20 shows off-line performances extracted 15 from Figs. 18 and 19.

Once again, the GA performance dominates random search for the duration of the experiment. It could be argued that the gap is not a large one but, as stated, the scores are normalized. The best network discovered 20 (by chance) after 60 generations took 67 epochs to learn the problem while the best network discovered by the GA learned the problem in 40 epochs. Further, it seems likely that we will be able to improve the performance of the GA through altered representation and better 25 parameter values, while there is no latitude for improvement in the performance of the random search

- 32 -

procedure. Finally, a caveat: we are running with a relatively small population, and our experiments have been limited to few generations-all of these results should therefore be interpreted with caution.

5

EXPERIMENT 2

Application: Digit Recognition

Optimization Criteria: Average fan-out and percent correct

10

Population Size: 30

Generations: 20

In this experiment, the criteria were the average fan-out and percentage correct, equally weighted (0.5). Learning rate was not given any direct influence

15

on the score. The percentage of correct digit identifications after training was determined by presenting each of the ten digits to the trained network and scoring a "Hit" if the output unit with maximal value corresponded to the correct digit. The average

20

fan-out is defined as the ratio of the number of weights to number of units; this metric is normalized and inverted, so that a large ratio of weights to units will detract from an individual's score. The question posed by this experiment is, can the system improve

25

performance by limiting fan-out? It is a potentially

- 33 -

interesting question to designers of neural network hardware, since high fan-outs are difficult to engineer in silicon. [Average fan-out is an approximation of an even more interesting quantity-maximal fan-out.] Our initial results are shown in Fig. 21.

The average fan-out in this experiment was $157/48 = 3.27$. This can be contrasted with the network shown for Experiment 1, which has an average fan-out that is almost an order of magnitude higher.

Learning was quite slow. In fact, the above network did not learn to within the error threshold that was prespecified as a termination criterion for training. (Learning to within the error threshold is not necessary to achieve perfect hit rates.) The connectivity structure of the network uses large receptive fields but low connection density. From a hardware implementation perspective, it would be better to optimize for small receptive fields and such an experiment is contemplated.

METRIC FOR LEARNING RATE

The metric chosen for learning rate requires some explanation. Because of limited computational resources, we cannot hope to train all networks until they achieve perfect accuracy on a given problem, or for that matter to any non-zero predetermined criterion. In

- 34 -

some cases, a network may require a hundred epochs while in others a million may be insufficient. Our compromise is to employ two criteria for halting the learning phase. Learning is halted under the first criterion

5 when rms error during the previous epoch was lower than a given threshold. The learning phase is terminated under the second criterion after a fixed number of epochs has been counted; this threshold is set by the experimenter according to the problem, but it is

10 typically between 100 to 5000 epochs. We nonetheless wish to compare all individuals on the same learning rate scale even though their training may have lasted different numbers of epochs and resulted in different final levels of accuracy. Our approximation is to

15 integrate the rms error curve over the learning phase for each individual. This "area under the learning curve" provides a rank that corresponds closely to our intuition about learning rate scales. Lower numbers imply better performance.

20

25

- 35 -

Appendix A

Syntax for Blueprint Representation in BNP:

```

<blueprint-spec> ::= <input-spec> <middle-spec> <output-spec>
<input-spec> ::= <area-spec> <projection-spec>
<middle-spec> ::= empty | <segment> | <middle-spec> <segment>
<output-spec> ::= <area-marker> <area-id> <dimension-subfield>
<area-spec> ::= <area-marker> <area-id> <dimension-subfield>
<learning-rule-spec>
<projection-spec> ::= <projection-marker> <projection-spec-field> |
<projection-spec>
<projection-marker> <projection-spec-field>
<area-spec> <projection-spec>
<segment> ::= empty
<area-marker> ::= empty
<area-id> ::= <binary-digits>
<dimension-subfield> ::= <total-size> <dim-spec> <dim-spec>
<dim-spec>
<learning-rule-spec> ::= <eta-initial-value> <slope-of-changing-eta>
<projection-marker> ::= empty
<projection-spec-field> ::= <radius-of-connectivity> <connection-density>
<target-address> <target-address-mode>
<learning-rule-spec>
<binary-digits> ::= <binary-digit> | <binary-digits> <binary-digit>
<binary-digit> <binary-digit> <binary-digit>
<binary-digit> <binary-digit> <binary-digit>
<binary-digit> <binary-digit> <binary-digit>
<binary-digit> <binary-digit> <binary-digit>
<radius-of-connection> <radius-of-connection>
<radius-of-connection>
<radius-of-connection> ::= <binary-digit> <binary-digit> <binary-digit>
<connection-density> ::= <binary-digit> <binary-digit> <binary-digit>
<target-address> ::= <binary-digit> <binary-digit> <binary-digit>
<target-address-mode> ::= <binary-digit>
<binary-digit> ::= 0 | 1

```

- 36 -

APPENDIX B

Backpropagation

5 Neural networks are constructed from two primitive elements: processing units and (directed) connections between units. The processing units are individually quite simple, but they are richly interconnected. Each connection typically has a real-valued weight associated with it, and this weight
10 indicates the effect the value of the unit at the source of the connection has on the unit at its destination. The output of a unit is some function of the weighted sum of its inputs:

15

$$o_j = f(\sum_i w_{ij} o_i - \theta_j) \quad (1)$$

20

Where o_j is the output of unit j , w_{ij} is the weight from unit i to unit j , and θ_j is the "threshold" or bias weight for unit j . The quantity $\sum_i w_{ij} o_i - \theta_j$ is usually referred to as the net input to unit j , symbolized net_j . The form of Eq. (1) that is usually employed with back-propagation is the sigmoid function:

25

$$f(x) = \frac{1}{1 + e^{-x}} \quad (2)$$

- 37 -

In most backpropagation networks, the units are arranged in layers and the networks are constrained to be acyclic. It can be shown that such "multi-layer feed-forward" networks can realize any mapping from a multi-dimensional continuous input space to a multi-dimensional continuous output space with arbitrarily high accuracy (Hecht-Nielsen, 1987; Lippmann, 1987; Lapedes & Farber, 1988).

5 The rule used to modify the weights is:
10

$$\Delta w_{ij} = \eta o_i \delta_j$$

This is the standard backpropagation learning rule. Here w_{ij} is the weight from unit i to unit j , o_i is the output of unit i , η is a constant that determines the learning rate, and δ_j is the error term for unit j . δ_j is defined differently for units in the output area and for units in "hidden" areas. For output units,

20

$$\delta_j = o_j'(t_j - o_j)$$

25 where o_j' is the derivative of o_j with respect to its net input (for the activation function of Eq. (2), this quantity is $o_j(1-o_j)$) and t_j is the target

- 38 -

value (the "desired output") for unit j. For hidden units, the target value is not known and the error term is computed from the error terms of the next "higher" layer:

5

$$\delta_j = o_j' \sum_k w_{jk} \delta_k$$

We have incorporated two extensions to most uses of backpropagation in our current implementation.

10 First, we use a recently discovered improvement of Eq. (3) (Samad, 1988):

$$\Delta w_{ij} = \eta (o_i + \delta_i) \delta_j$$

15 This equation uses the anticipated value of the source unit of a weight instead of the current computed value. In some cases, orders of magnitude faster learning is achieved.

20 Second, we allow the value of η to decrease as learning proceeds. That is, η is now a variable, and the learning rule actually used is:

$$\Delta_{ij} = \eta_t (o_i + \delta_i) \delta_j$$

25

- 39 -

where η_t is the value of η at the t th iteration through the training set. At the end of each iteration, η is changed according to the following formula:

5

$$\eta_{t+1} = \eta \text{ slope } \eta_t$$

where η slope is a parameter that determines the rate of decay of η . It has been experimentally observed that using a high value of η initially and then gradually decreasing it results in significantly faster learning than using a constant η . Both η slope and the initial value of η (η_0) are given by the projection specification in the blueprint.

10

15

20

25

- 40 -

Appendix C

References

5 Baker, J.E. (1987). Reducing bias and inefficiency in the selection algorithm. *Proceedings of the Second International Conference on Genetic Algorithms*, pp. 14-21.

10 Bethke, A.D. (1980). *Genetic Algorithms As Function Optimizers*. Doctoral dissertation, University Microfilms International: Ann Arbor, MI.

15 Booker, L. (1987). Improving search in genetic algorithms. *Genetic Algorithms and Simulated Annealing*. L. Davis (Ed.) Morgan Kaufmann, Los Altos, CA.

20 Caruana, R.A., and J.D. Schaffer. (1988). Representation and hidden bias: gray vs. binary coding for genetic algorithms. *Proceedings of the 5th International Conference on Machine Learning*, Ann Arbor, MI, pp.153-161.

25 Carpenter, G.A., and S. Grossberg. (1988). The ART of adaptive pattern recognition by a self-organizing neural network. *Computer*. Vol. 21, No. 3, pp. 77-88.

30 DeJong, K. (1980). Adaptive systems design: a genetic approach. *IEEE Transactions on Systems, Man and Cybernetics*, Vol. 10, No. 9, pp. 566 - 574.

35 Dolan, C.P., and M. G. Dyer. (1987). Towards the evolution of symbols. *Proceedings of the Second International Conference on Genetic Algorithms*, pp. 123 - 131.

40 Dress, W.B., and J. R. Knisley. (1987). A Darwinian approach to artificial neural systems. *Proceedings of the IEEE Conference on Systems, Man and Cybernetics*, pp. 572 - 577.

45 Edelman, G. (1988). *Neural Darwinism*. 1988. Harper and Row.

50 Goldberg, D.E. (1988). *Genetic Algorithms in Search, Optimization & Machine Learning*. Addison-Wesley: Reading, MA.

55 Goldberg, D.E. (1987). Simple genetic algorithms and the minimal, deceptive problem. *Genetic Algorithms and Simulated Annealing*. L. Davis, Ed. Morgan Kaufmann, Los Altos, CA.

60 Hinton, G.E. (1988). Generative back-propagation. Paper presented at the International Neural Networks Society Meeting.

65 Huang, W.Y., and R.P. Lippmann. (1987). Neural net and traditional classifiers. *Neural Information Processing Systems*. D.Z. Anderson Ed. American Institute of Physics.

70 Holland, J. (1986). *Escaping Brittleness: The Possibilities of General-Purpose Learning Algorithms Applied to Parallel Rule-based Systems*. *Parallel Learning*, Vol. 2, Morgan Kaufmann, Los Altos, CA.

Holland, J. (1973). *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor.

Kohonen, T.E. (1984). *Self-Organization and Associative Memory*. Springer-Verlag.

5 Linsker, R. (1988). Self-organization in a perceptual network. *Computer*, Vol. 21, No. 3, pp. 105-117

Le Cun, Y. (1986). Learning process in an asymmetric threshold network. *Disordered Systems and Biological Organization*. E. Bienenstock et al. Ed. Springer Verlag.

10 Parker, D.B. (1985). *Learning-Logic*. TR-47, Center for Computational Research in Economics and Management Science, MIT, April.

Rumelhart, D.E. (1988). Plenary talk at the IEEE International Conference on Neural Networks, San Diego.

15 Rumelhart, D.E., G.E. Hinton, and R.J. Williams. (1985). *Learning Internal Representations by Error Propagation*. ICS Report 8506, Institute for Cognitive Science, University of California, San Diego.

Samad, T. (1988). Back-propagation is significantly faster with anticipated values of source units. *Neural Networks*, Vol. 1, Sup. 1. (Abstracts of the First Annual INNS Meeting.)

20 Werbos, P.J. (1974). *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences*. Ph.D. thesis, Department of Applied Mathematics, Harvard University, Cambridge, Mass.

- 42 -

It is claimed:

1. A method for synthesizing designs for neural networks which involves the use of a selected learning algorithm and a particular subject to be learned, comprising the steps of:
 - 5 A. devising a bit string parametric representation of a neural network architecture having relevant parameters,
 - B. generating a first generation of network 10 blueprints based on said representation which jointly include a range of values for each of said parameters,
 - C. generating respective neural network 15 architectures based on the current generation of said blueprints,
 - D. training each of said network architectures presently defined in step C via said selected learning algorithm and said subject matter,
 - E. testing each of said network architectures presently defined in step C with test patterns 20 corresponding to said subject matter for testing the receptiveness of each of said network architectures presently defined in step C to the affect of said training.
 - F. performing an evaluation for each of said 25 network models presently defined in step C after said testing thereof relative to performance and cost factors of interest and assigning a score thereto representing the results of said evaluation,

- 43 -

5 G. selecting candidates from said network architectures presently identified in step C in accordance with some rationale and applying at least one operator thereto to produce a new generation of network blueprints which shall be identified as the current generation of network blueprints based on said representation, and

H. returning to step C and continuing the process.

10

2. A method according to claim 1 wherein said operator is a genetic operator.

15 3. A method for synthesizing designs for neural networks each of which comprise,

20 a plurality of computational units, a plurality of hierarchically arranged layer areas including input and output layer areas and zero or more hidden layer areas therebetween, each of said layer areas being defined by a number of said units, connecting means connecting source groups of said units in said layer areas other than said output layer area with object groups of said units in said layer areas other than said input layer area, said connecting means being grouped in sets deemed projections with each of said projections extending from one of said layer areas to another of said layer areas,

25

- 44 -

 said method comprising the steps of:

 A. providing a substring format for specifying
 each of said layer areas with said format having one
 first type part deemed a layer area parameter specifier
 5 and at least one second type part for each of said
 projections deemed a projection specifier,

 said first type part comprising a layer area
 identifying address section, a total size section
 denoting the corresponding number of said units thereof,
 10 and a dimension section denoting the configuration
 formed by said units,

 each said second type part being dedicated to
 one of said projections deemed a subject projection,
 said second part type comprising a target address
 15 section for identifying one of said layer areas deemed a
 target layer area to which said subject projection is
 directed, a mode of address section for said subject
 projection, a dimension section for denoting the
 configuration of an object field for said subject
 20 projection in said target layer area, a connection
 density section for denoting the connectivity of said
 subject projection to said object field, and at least
 one learning rule parameter section.

 B. devising a bit string parametric
 25 representation of a neural network architecture based on
 said substring format and having relevant parameters,

- 45 -

C. generating a first generation of network blueprints based on said representation which jointly include a range of values for each of said parameters,

5 D. generating respective neural network architectures based on the current generation of said blueprints,

E. training each of said network architectures presently defined in step D via said selected learning algorithm and said subject matter,

10 F. testing each of said network architectures presently defined in step D with test patterns corresponding to said subject matter for testing the receptiveness of each of said network models presently defined in step D to the affect of said training,

15 G. performing an evaluation for each of said network architectures presently defined in step D after said testing thereof relative to performance and cost factors of interest and assigning a score thereto representing the results of said evaluation.

20 H. selecting candidates from said network architectures presently identified in step D in accordance with some rationale and applying at least one genetic operator thereto to produce a new generation of network blueprints which shall be identified as the
25 current generation of network blueprints based on said representation, and

- 46 -

I. returning to step D and continuing the process.

5

10

15

20

25

1 / 14

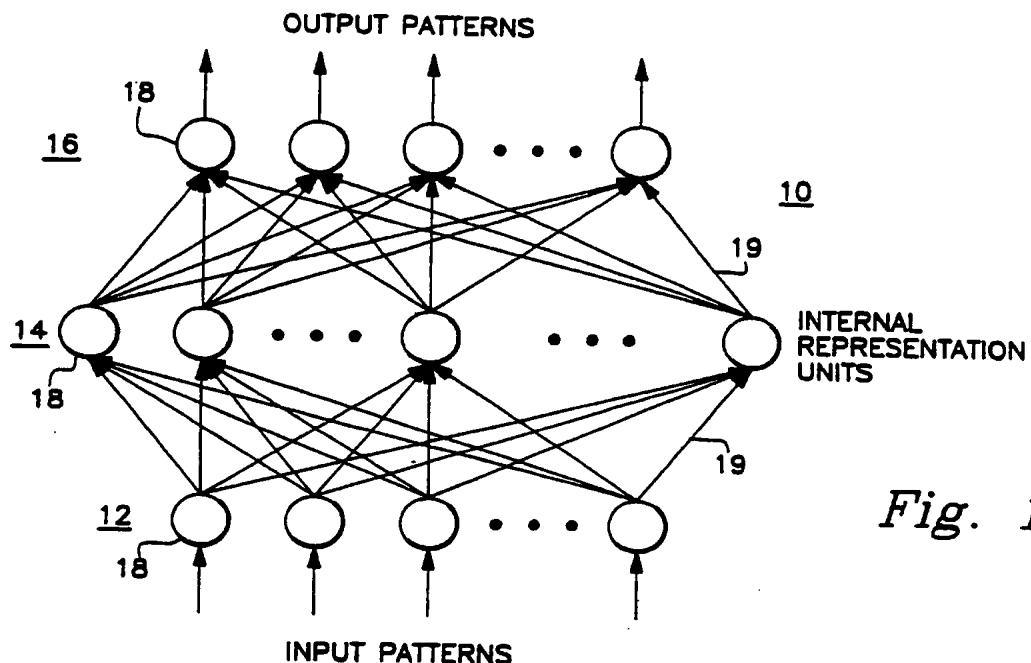


Fig. 1

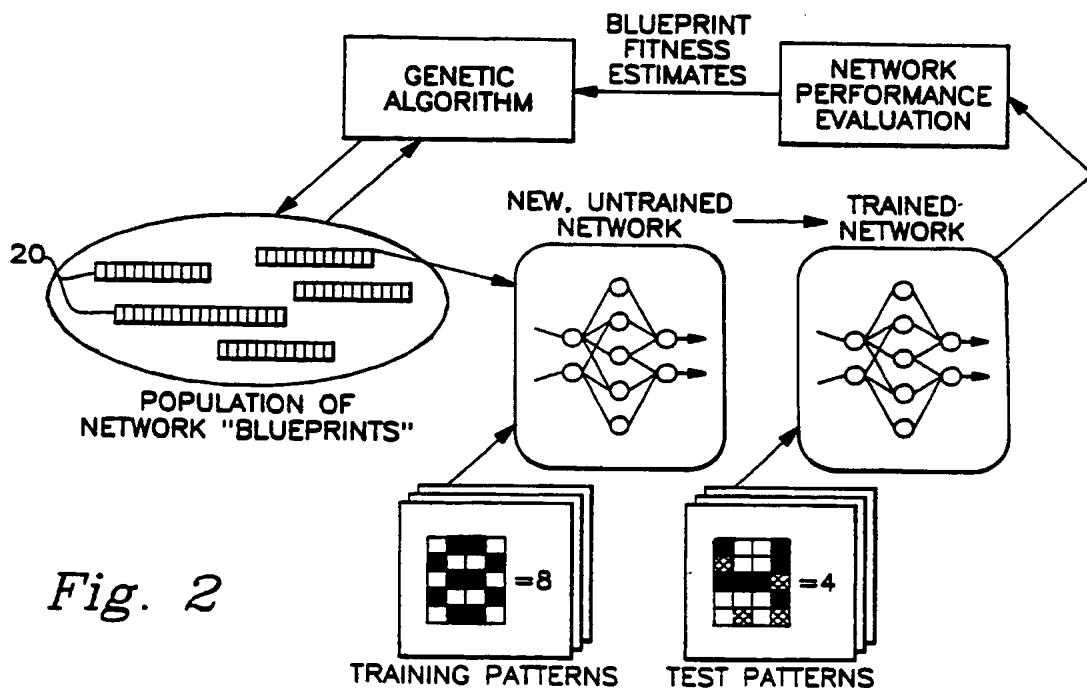


Fig. 2

SUBSTITUTE SHEET

2 / 14

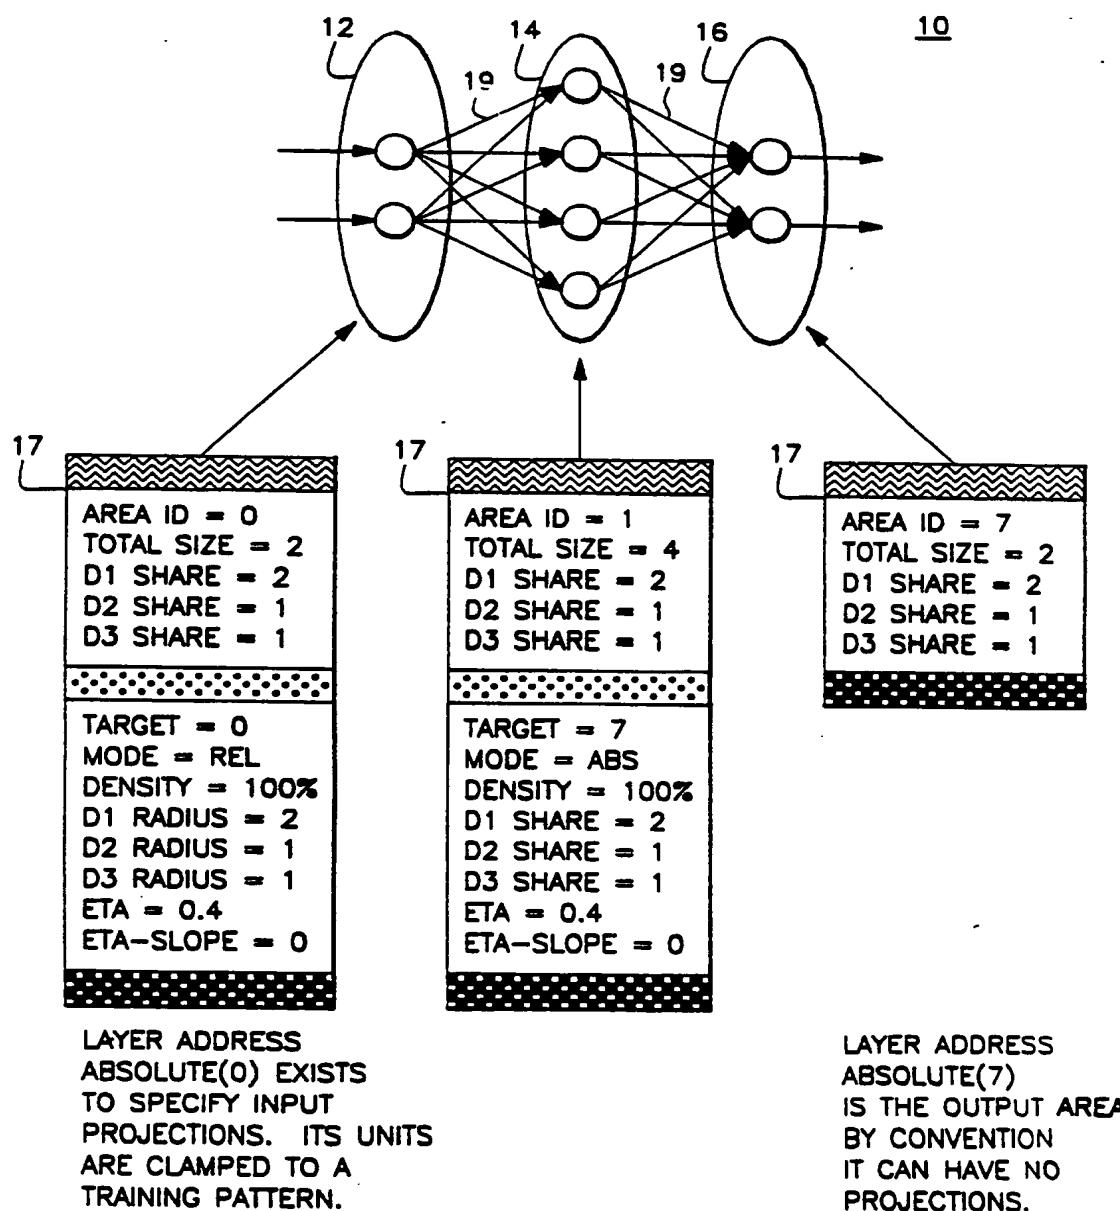


Fig. 3

SUBSTITUTE SHEET

3 / 14

AREA SPECIFICATION SUBSTRING

■■■■■ START OF AREA MARKER

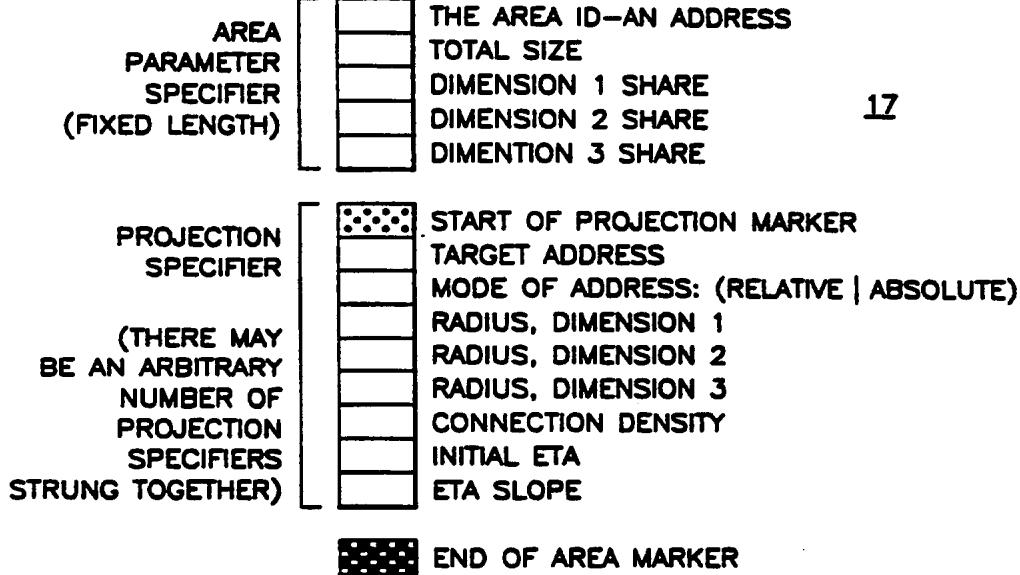


Fig. 4

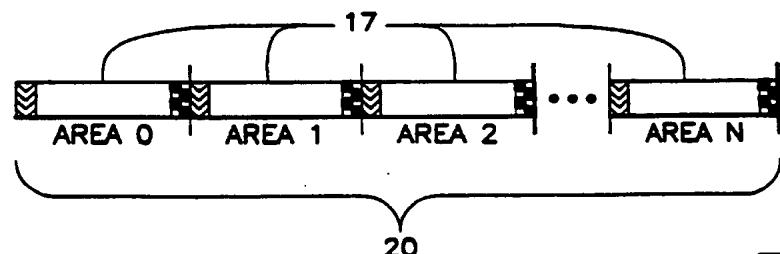


Fig. 5

SUBSTITUTE SHEET

4 / 14

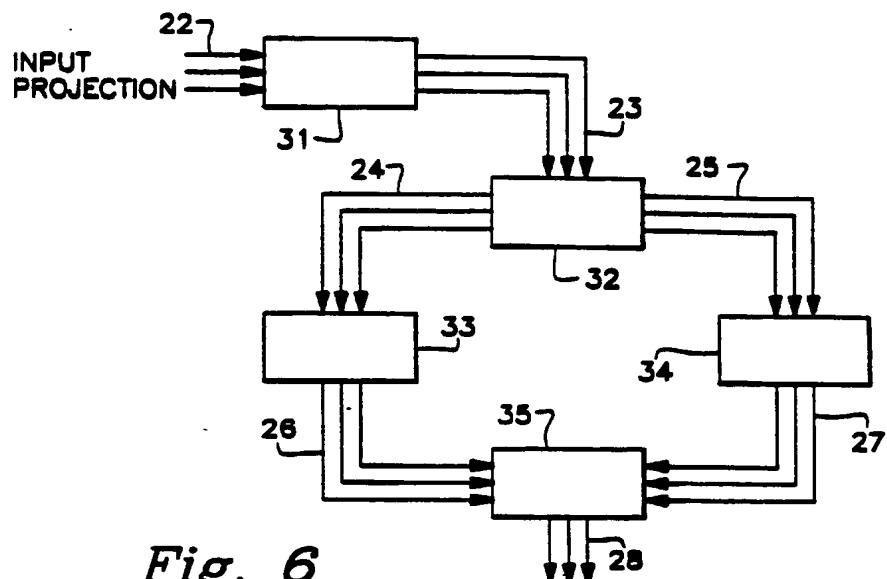


Fig. 6

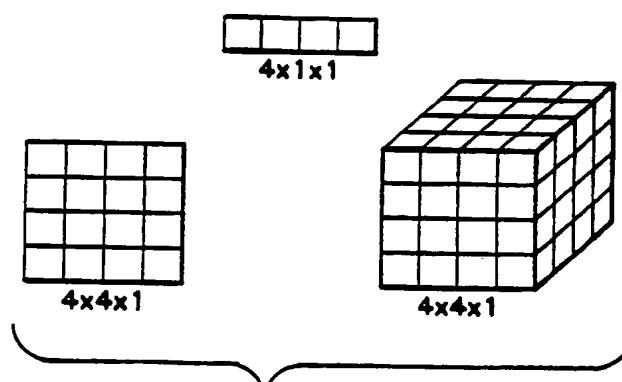
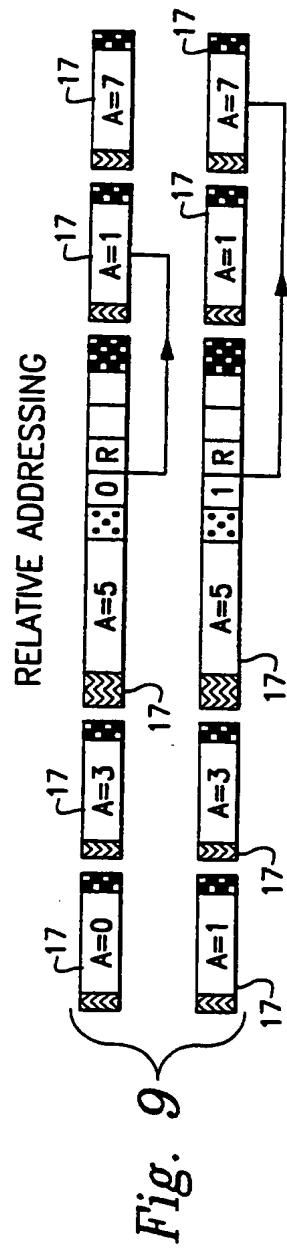
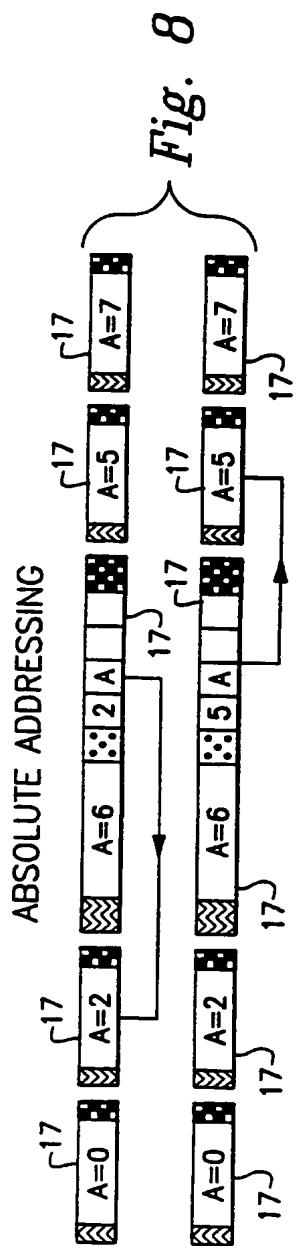
OUTPUT
PROJECTION

Fig. 7

SPATIAL ORGANIZATION OF AREAS.
EACH SMALL SQUARE REPRESENTS A UNIT.
THE LABELS INDICATE THE "SHARE" IN EACH DIMENSION.

SUBSTITUTE SHEET



6 / 14

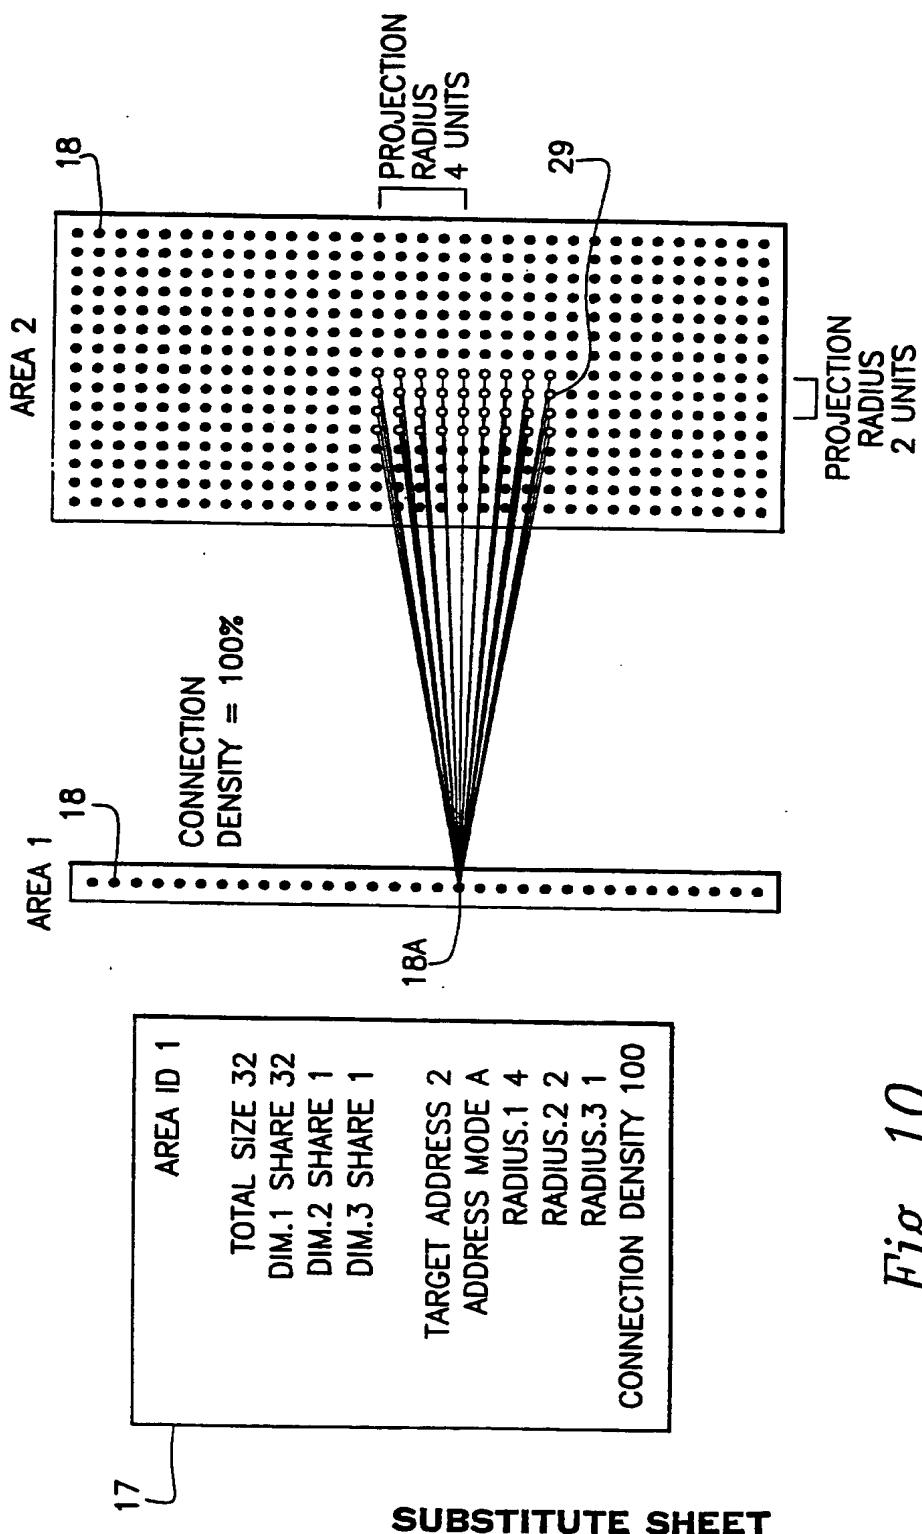


Fig. 10

7 / 14

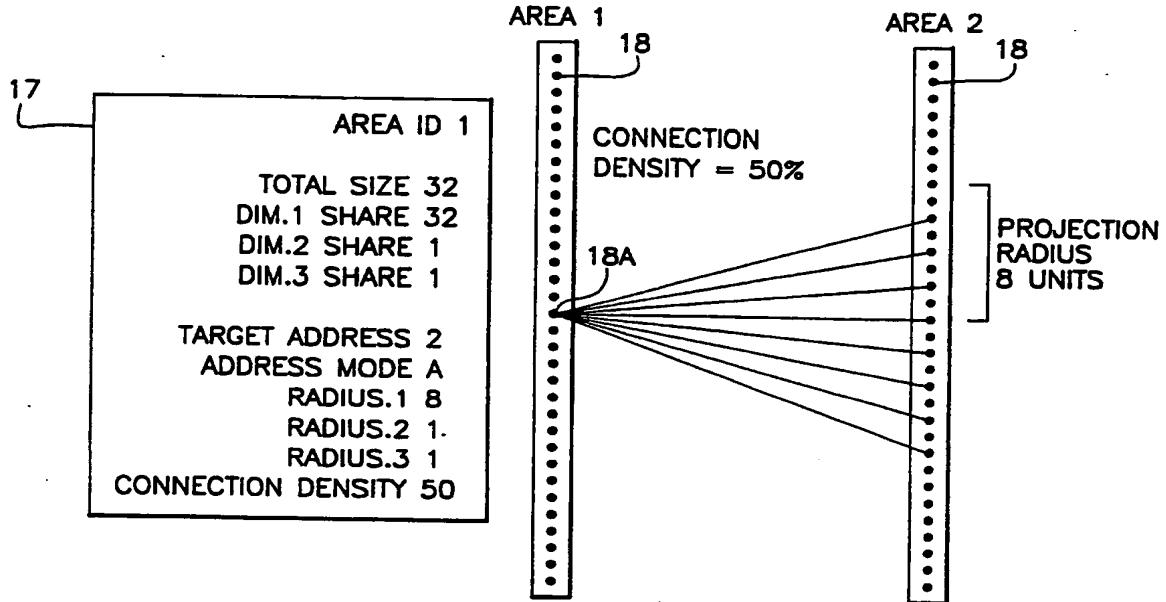


Fig. 11

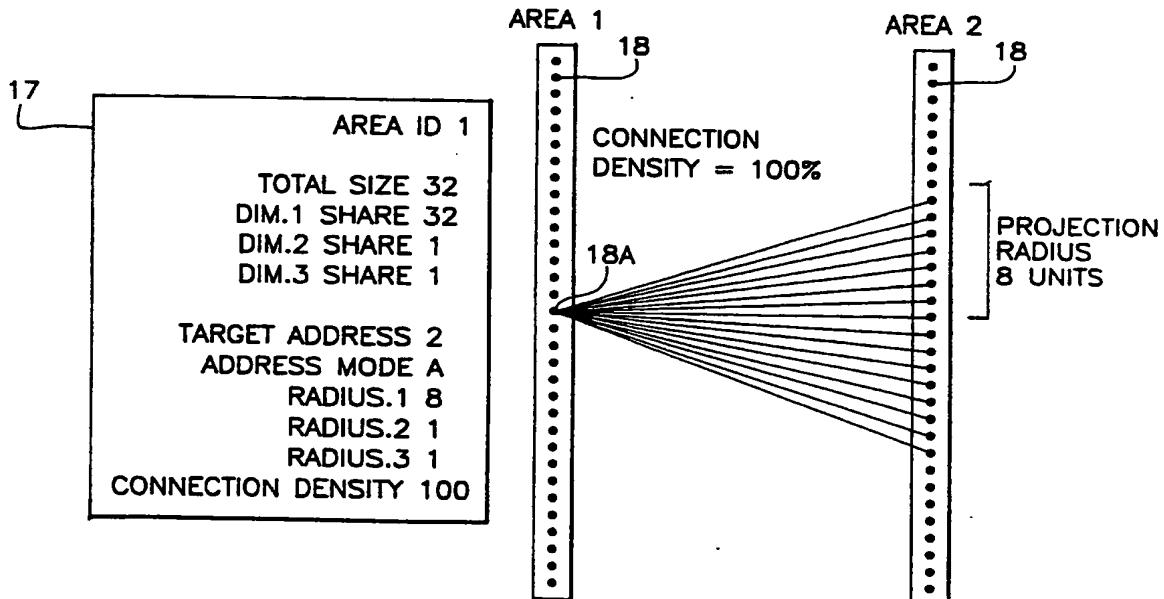
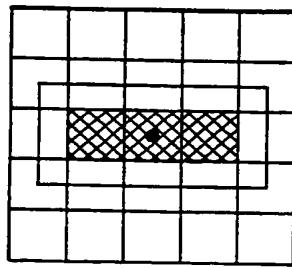


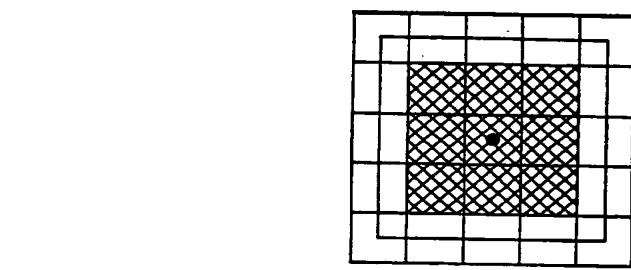
Fig. 12

SUBSTITUTE SHEET



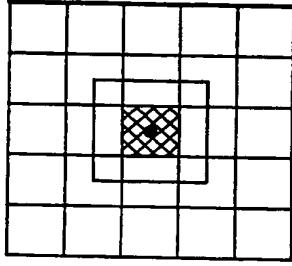
RADI ARE 1x2, SO A 2x4 BOX SURROUNDS THE CENTER DOT IT COMPLETELY ENCLOSES 3 UNITS, SHOWN SHAD

Fig. 13 A



RADI ARE 2x2, SO A 4x4 BOX SURROUNDS THE CENTER DOT IT COMPLETELY ENCLOSES 9 UNITS, SHOWN SHAD

Fig. 13 B



RADI ARE 1x1, SO A 2x2 BOX SURROUNDS THE CENTER DOT IT COMPLETELY ENCLOSES 1 UNIT, SHOWN SHAD

Fig. 13 C

Fig. 14 A



Fig. 14 B

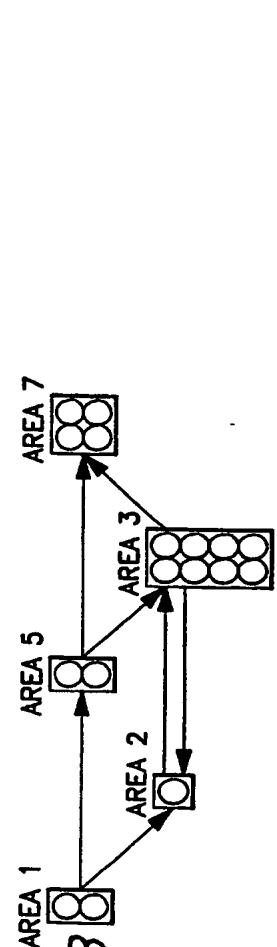
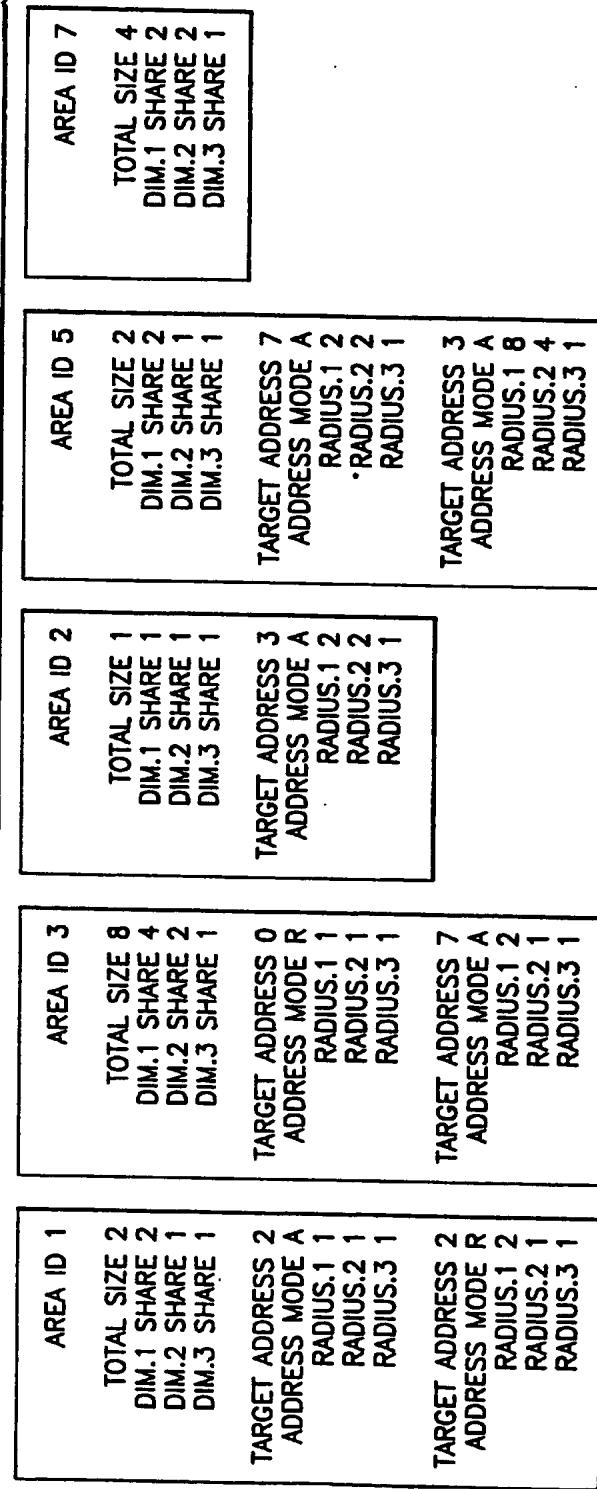


Fig. 14 C



10 / 14

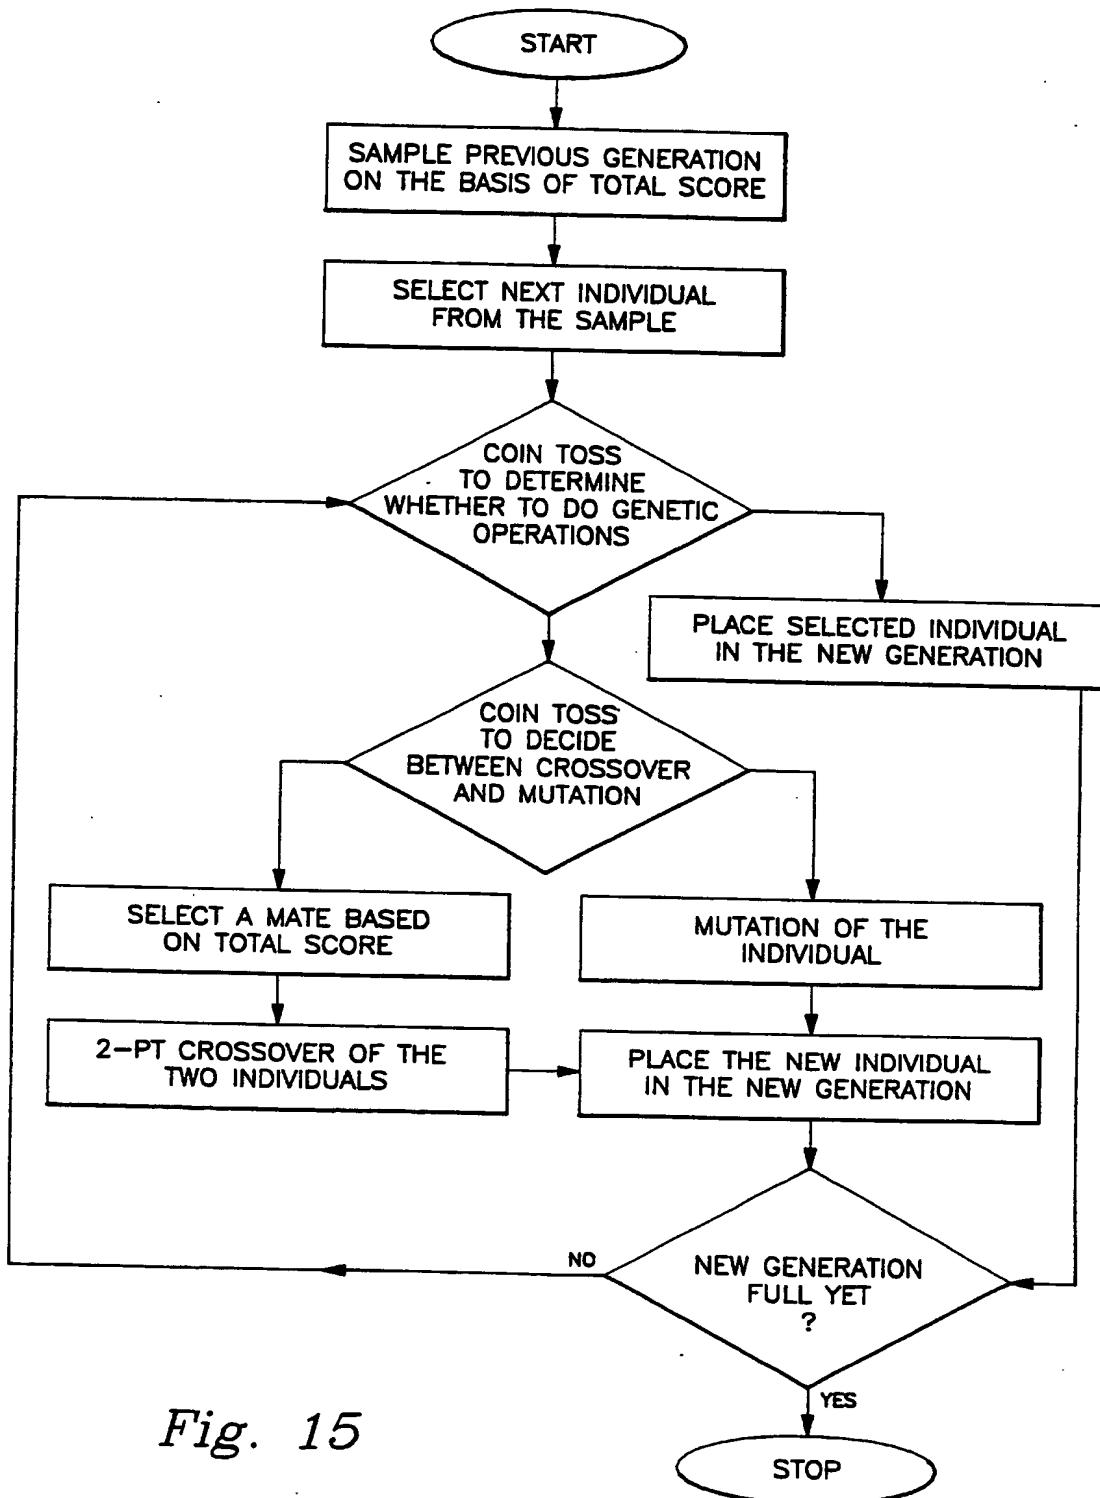
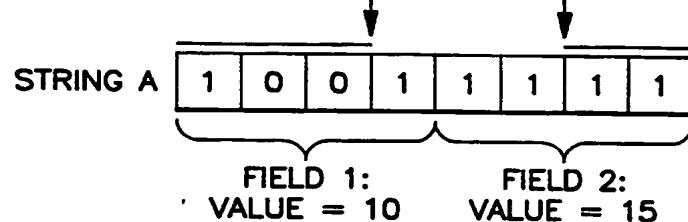
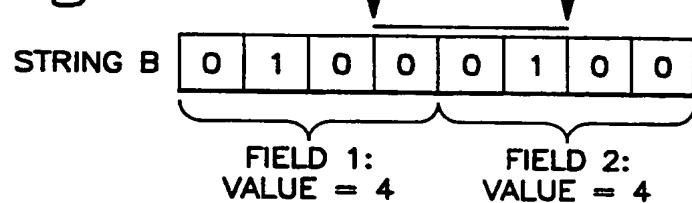
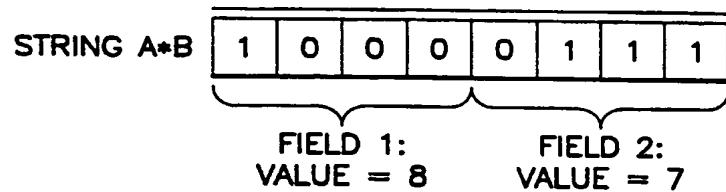
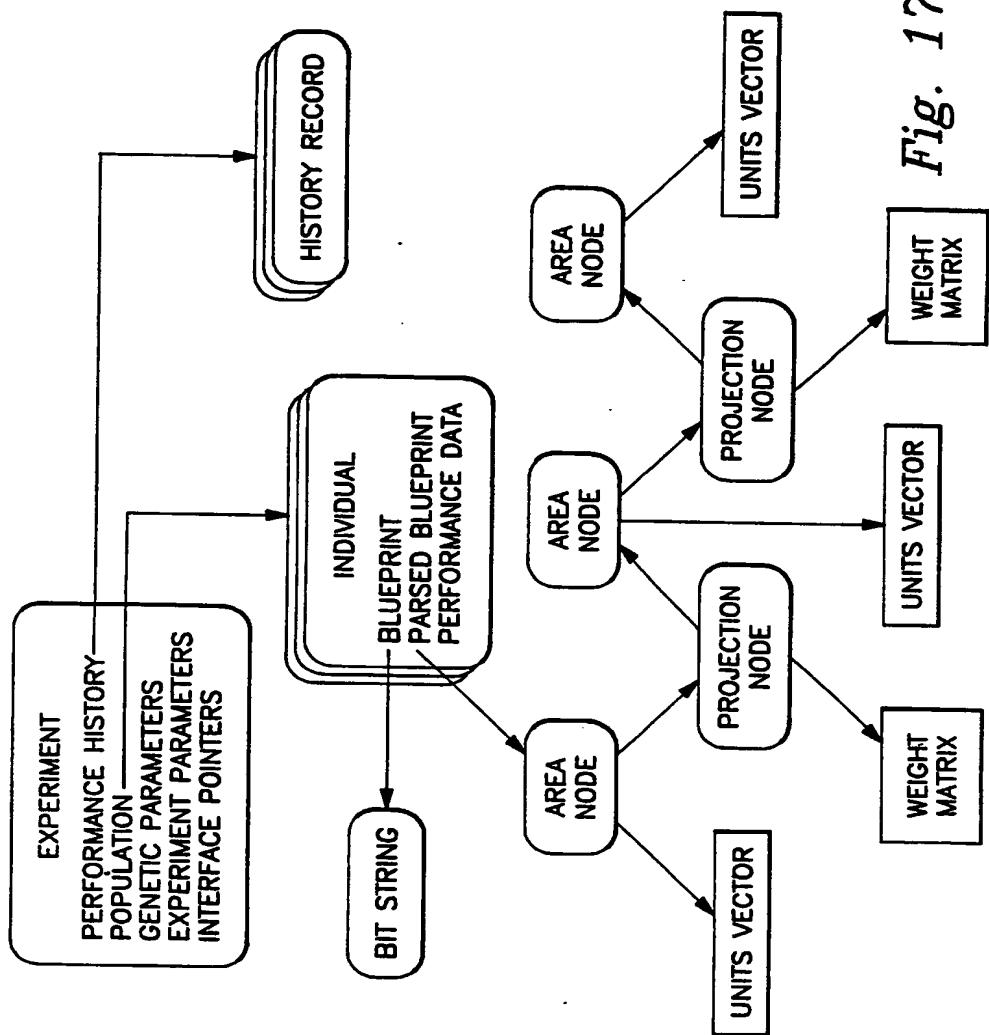


Fig. 15

SUBSTITUTE SHEET

11 / 14

Fig. 16A (RANDOMLY SELECTED) CROSSOVER POINTS*Fig. 16B* CROSSOVER POINTS*Fig. 16C* NEW STRING



13 / 14

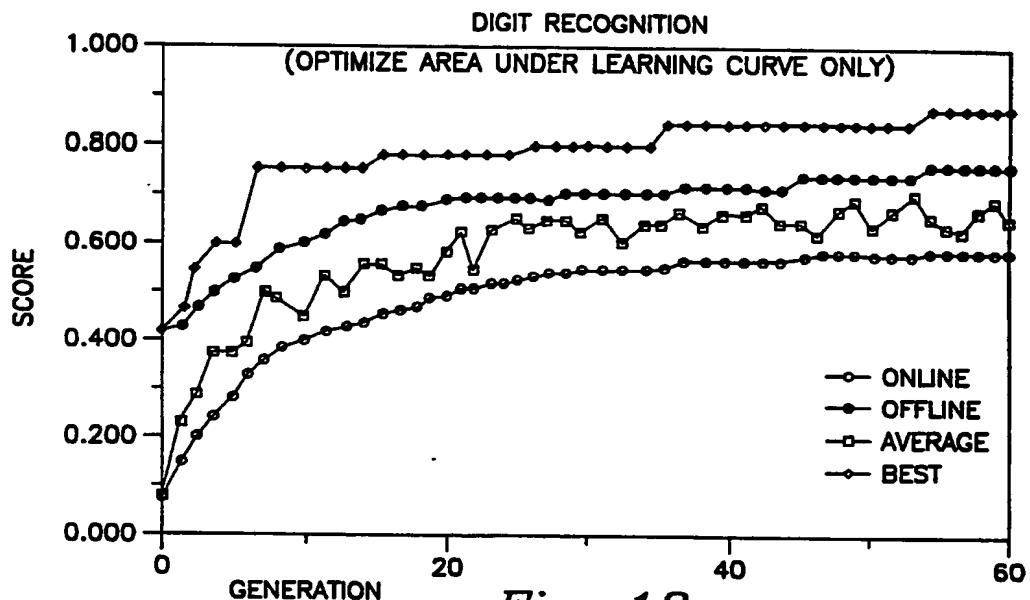


Fig. 18

PERFORMANCE IN THE DIGIT RECOGNITION PROBLEM. SCORES DEPENDED ONLY ON THE RATE OF LEARNING (UNDER THE LEARNING CURVE)

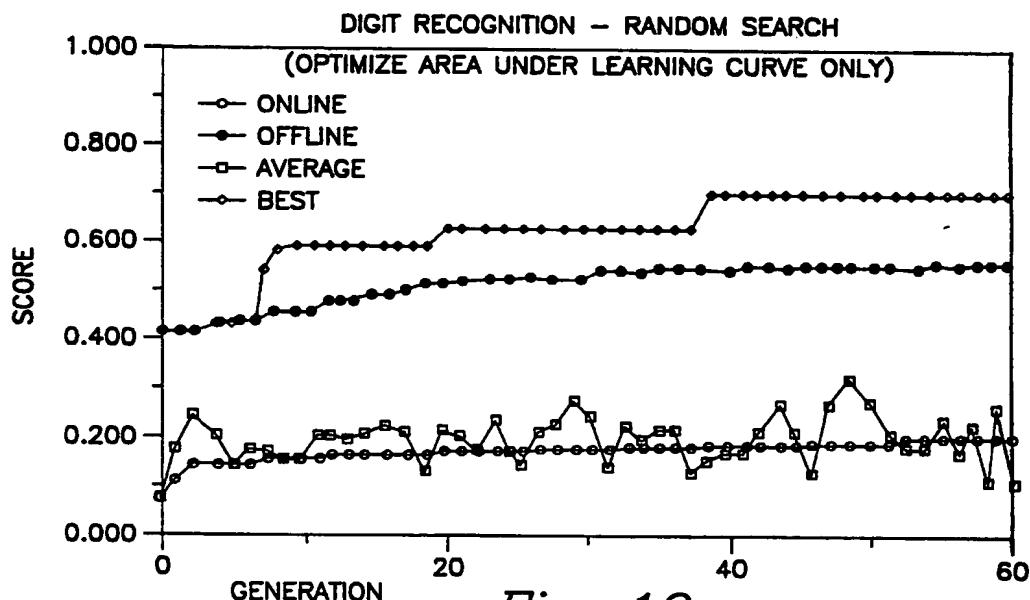


Fig. 19

PERFORMANCE ON THE DIGIT RECOGNITION PROBLEM WITH RANDOM SEARCH REPLACING THE GENETIC OPERATORS. ONLY THE AREA UNDER THE LEARNING CURVE CONTRIBUTED TO THE SCORE

SUBSTITUTE SHEET

14 / 14

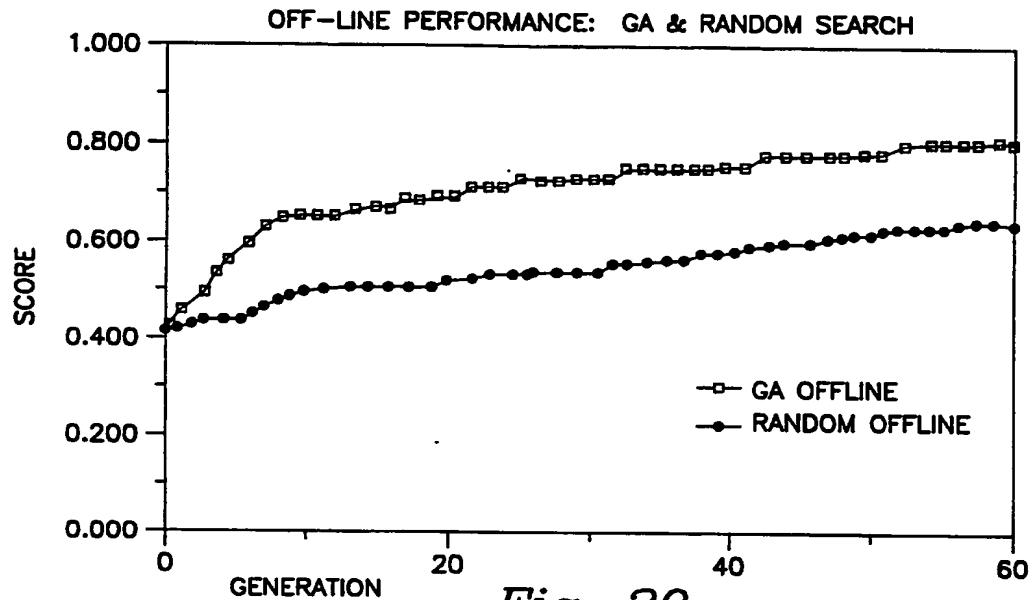


Fig. 20

OFF-LINE PERFORMANCE ON THE DIGIT RECOGNITION PROBLEM,
COMPARING THE GA AND RANDOM SEARCH.

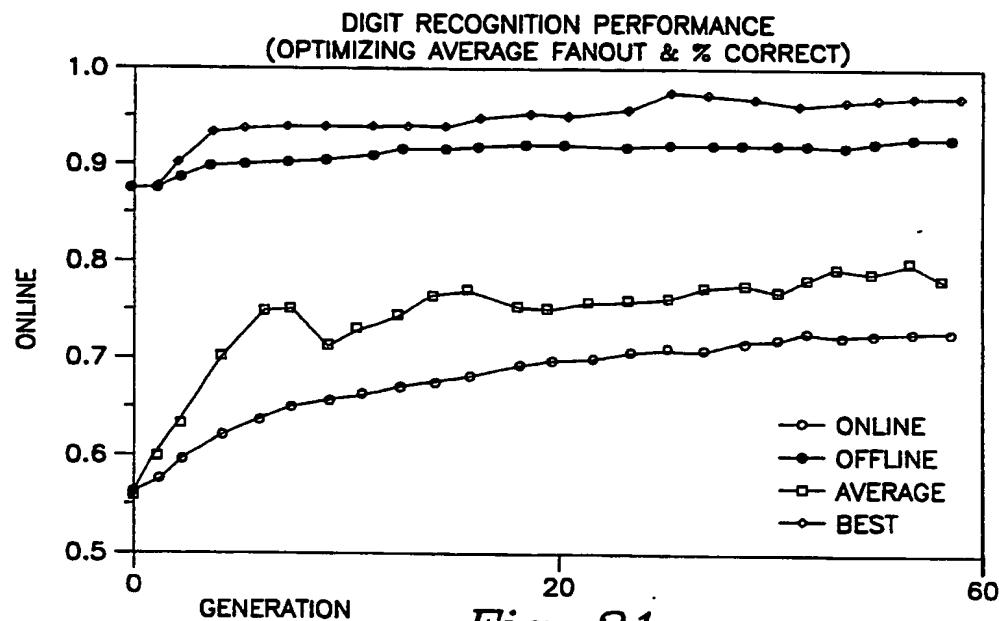


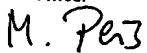
Fig. 21

PERFORMANCE CURVES FOR THE DIGIT RECOGNITION PROBLEM WITH EQUAL
WEIGHT ON THE PERCENT CORRECT AND THE AVERAGE FAN-OUT

SUBSTITUTE SHEET

INTERNATIONAL SEARCH REPORT

International Application No. PCT/US 90/00828

I. CLASSIFICATION OF SUBJECT MATTER (If several classification symbols apply, indicate all) ⁶ According to International Patent Classification (IPC) or to both National Classification and IPC IPC5: G 06 F 15/18													
II. FIELDS SEARCHED <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th colspan="2" style="text-align: center; padding: 2px;">Minimum Documentation Searched⁷</th> </tr> <tr> <th style="text-align: center; padding: 2px;">Classification System</th> <th style="text-align: center; padding: 2px;">Classification Symbols</th> </tr> </thead> <tbody> <tr> <td style="text-align: center; padding: 2px;">IPC5</td> <td style="text-align: center; padding: 2px;">G 06 F</td> </tr> <tr> <td colspan="2" style="text-align: center; padding: 2px;">Documentation Searched other than Minimum Documentation to the Extent that such Documents are Included in Fields Searched⁸</td> </tr> </tbody> </table>		Minimum Documentation Searched ⁷		Classification System	Classification Symbols	IPC5	G 06 F	Documentation Searched other than Minimum Documentation to the Extent that such Documents are Included in Fields Searched ⁸					
Minimum Documentation Searched ⁷													
Classification System	Classification Symbols												
IPC5	G 06 F												
Documentation Searched other than Minimum Documentation to the Extent that such Documents are Included in Fields Searched ⁸													
III. DOCUMENTS CONSIDERED TO BE RELEVANT⁹ <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: center; padding: 2px;">Category¹⁰</th> <th style="text-align: center; padding: 2px;">Citation of Document,¹¹ with indication, where appropriate, of the relevant passages¹²</th> <th style="text-align: center; padding: 2px;">Relevant to Claim No.¹³</th> </tr> </thead> <tbody> <tr> <td style="text-align: center; padding: 2px;">A</td> <td style="text-align: center; padding: 2px;">US, A, 3273125 (C.V.JAKOWATZ) 13 September 1966, see the whole document</td> <td style="text-align: center; padding: 2px;">1-3</td> </tr> <tr> <td style="text-align: center; padding: 2px;">A</td> <td style="text-align: center; padding: 2px;">US, A, 4591980 (HUBERMAN ET AL) 27 May 1986, see the whole document</td> <td style="text-align: center; padding: 2px;">1-3</td> </tr> <tr> <td style="text-align: center; padding: 2px;">A,P</td> <td style="text-align: center; padding: 2px;">EP, A2, 325119 (NESTOR INC) 26 July 1989, see the whole document</td> <td style="text-align: center; padding: 2px;">1-3</td> </tr> </tbody> </table>		Category ¹⁰	Citation of Document, ¹¹ with indication, where appropriate, of the relevant passages ¹²	Relevant to Claim No. ¹³	A	US, A, 3273125 (C.V.JAKOWATZ) 13 September 1966, see the whole document	1-3	A	US, A, 4591980 (HUBERMAN ET AL) 27 May 1986, see the whole document	1-3	A,P	EP, A2, 325119 (NESTOR INC) 26 July 1989, see the whole document	1-3
Category ¹⁰	Citation of Document, ¹¹ with indication, where appropriate, of the relevant passages ¹²	Relevant to Claim No. ¹³											
A	US, A, 3273125 (C.V.JAKOWATZ) 13 September 1966, see the whole document	1-3											
A	US, A, 4591980 (HUBERMAN ET AL) 27 May 1986, see the whole document	1-3											
A,P	EP, A2, 325119 (NESTOR INC) 26 July 1989, see the whole document	1-3											
* Special categories of cited documents: ¹⁰ "A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier document but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed *T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention *X document of particular relevance, the claimed invention cannot be considered novel or cannot be considered to involve an inventive step *Y document of particular relevance, the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art *Z document member of the same patent family													
IV. CERTIFICATION													
Date of the Actual Completion of the International Search	Date of Mailing of this International Search Report												
15th June 1990	09.07.90												
International Searching Authority	Signature of Authorized Officer												
EUROPEAN PATENT OFFICE	 M. PEIS												

ANNEX TO THE INTERNATIONAL SEARCH REPORT
ON INTERNATIONAL PATENT APPLICATION NO.PCT/US 90/00828

SA 35361

This annex lists the patent family members relating to the patent documents cited in the above-mentioned international search report.
The members are as contained in the European Patent Office EDP file on 24/05/90
The European Patent office is in no way liable for these particulars which are merely given for the purpose of information.

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US-A- 3273125	13/09/66	NL-A- 269964	00/00/00
US-A- 4591980	27/05/86	DE-A- 3501272 GB-A-B- 2154343 JP-A- 60183645	29/08/85 04/09/85 19/09/85
EP-A2- 325119	26/07/89	NONE	

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

EPO FORM P0479